A high stability ka-band radiometer for tropospheric water vapor measurements

Alan B. Tanner
Jet Propulsion Laboratory, California Inst. of Technology
4800 Oak Grove Dr.
Pasadena, CA 91109
818-354-7477
alan.b.tanner@jpl.nasa.gov

2001 IEEE Aerospace Conference
DRAFT 11/20/00 ID=155

Abstract- The design of two new high stability microwave water vapor radiometers is presented along with a performance evaluation. The radiometers operate next to a spacecraft tracking station at NASA's Goldstone facility in California where they will be used to calibrate tropospheric path delay fluctuations during an upcoming gravity wave search experiment (GWE) involving the Cassini spacecraft. Observing frequencies of the radiometers are 22.2 GHz, 23.8 GHz, and 31.4GHz, and the antenna beamwidth is 1 degree. The instruments are room temperature Dicke radiometers with additive noise injection for gain calibration. Design highlights include: (1) a practical temperature control system capable of stabilizing the entire receiver to a few millikelvin from day to night; (2) multiply redundant noise diode injection circuits with 30 parts per million RF power stability; and (3) a voice coil actuated waveguide vane attenuator which is used as a high performance Dicke switch. Performance of the radiometers is evaluated from inter-comparisons of the two radiometers and from continuous tip-curve calibrations spanning nearly one year. Structure function analysis of the inter-comparison data indicates that the brightness temperature stability of these radiometers is better than 0.01 Kelvin on 1,000 to 10,000 second time scales. Analysis of tip-curve calibrations indicate RMS errors of approximately 0.05 K on 30-day time scales, and 0.15K on one year time scales.

TABLE OF CONTENTS

1. Introduction

Water vapor adds a variable delay to radio signals passing through the earth's troposphere that can limit the phase stability and coherence time for VLBI or for experiments that involve precise spacecraft tracking. One technique to correct these errors involves measuring the water vapor along a line of sight with a water vapor radiometer (WVR)[1]. WVR's measure the thermal radio emissions of water vapor near a resonance at 22 GHz, and the radio path delay is inferred from these data using models and estimates of the atmospheric temperature and pressure. There are a variety of errors associated with this technique, including model errors and uncertainty in the vertical distribution of temperature and water, but past experiments have been largely limited by the accuracy and stability of the WVR itself [2],[3].

At JPL, an upcoming series of gravitational wave search experiments (GWE) involving the Cassini spacecraft will require the most sensitive measure of the water vapor path delay to date. Starting in December of 2001, and during the spacecraft's cruise to Saturn, the Cassini GWE will be conducted from a specially equipped ground station at NASA's Goldstone Deep Space Communication Complex (GDSCC) in California. The principle GWE measurements will be of very small perturbations- potentially less than ~ 1mm- in the spacecraft tracking data on time scales of 1,000 to 10,000 seconds. Of the various errors associated with this experiment- including those of the time standards, mechanical stability of the ground station and spacecraft, etc.- it is the tropospheric path delay corrections which will likely determine the detection threshold for the GWE [4]. When translated into radiometric specifications, Cassini GWE goals call for a WVR that is stable to better than 0.01K in brightness temperature. This represents about a ten-fold improvement over the demonstrated performance of any previous WVR, and has motivated the development of the radiometers described in this paper.

This paper presents the design and performance evaluation of two WVR's which were developed for the Cassini GWE. Dubbed the 'advanced' water vapor radiometers (AWVR's), these instruments are based on laboratory tests and prototype work which was previously reported by [5]. The laboratory work produced a number of innovations which will be briefly reviewed in section 3. Section 4 will briefly describe the internal calibration algorithm, and Section 5 will present a performance evaluation the radiometers based on side by side comparisons and from long term calibration data spanning nearly one year.

2. System Overview

The two AWVR units, designated by serial numbers A-1 and A-2, are located approximately 40 meters to the south-east and south-west, respectively, of the much larger tracking antenna, designated DSS-25, where the Cassini GWE will be conducted. During a spacecraft track both AWVR's will be pointed in the same direction as the DSS-25 antenna. During off-hours the AWVR's perform continuous sweeps in elevation as part of a 'tip-curve' procedure that is the basis for the system calibration. Two identical radiometers have been built to ensure reliability and to provide a cross-check of the system performance.

Figure 1 shows one of the AWVR's. The AWVR consists of the radiometer 'box' in the upper left, the supporting pedestal, and the offset parabolic reflector. The overall height of the system is about 3 meters, and the reflector size is about 1.6 by 1.9 meters. The radiometer 'box' contains the antenna feedhorn (hidden behind the circular radome on the face of the enclosure), RF electronics, data system, and temperature control systems. The pedestal includes the azimuth positioner at the base of the structure and the elevation positioner which rotates the reflector on a focal axis that is inclined by 35 degrees from horizontal. This 35 degree inclination is complementary to a 55 degree offset between the focal axis and the main beam of the antenna. Figure 1 shows the antenna as it points the main beam to zenith. Rotation of the elevation positioner to other angles results in a conical sweep of the antenna beam about the focal axis, and requires a coordinate transformation to convert between positioner angles and the 'true' azimuth and elevation, as referenced to the local horizon. In practice these conversions are calculated by the controlling computer and are transparent to the user.

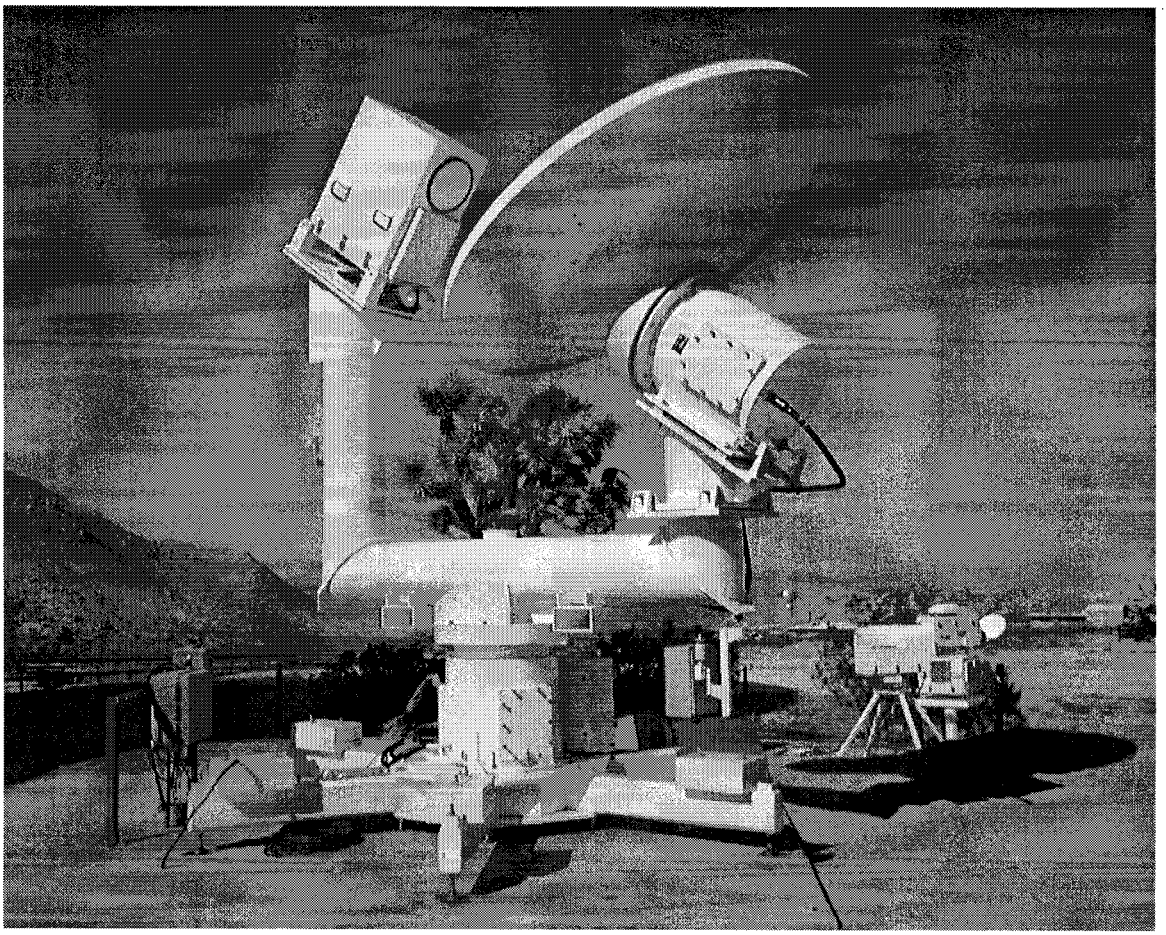

Figure 1- One of two AWVR's in operation at Goldstone, CA.

A short list of specification for the AWVR are summarized in Table 1. The frequency selections of Table 1 are base on work by [2] to minimize the uncertainty in the path delay inversion algorithm. The antenna beamwidth is designed to minimize the mismatch with the 34 meter diameter tracking antenna at a range of 2 km, which is approximately the midpoint of the troposphere (note that a 1 degree beam projects a 34 meter spot diameter at 2 km range, which roughly matches the near-field pattern of the DSS-25 tracking antenna) [6], [7]. The pointing requirements of Table 1 were derived from the sensitivity of the measurement to air mass- which increases rapidly (as the secant of elevation angle) as the antenna beam approaches the horizon.

One specification missing in Table 1 is the beam efficiency. This often quoted figure of merit for a radiometer antenna is difficult to apply in this case since it does not distinguish between near and far sidelobes. Without any such distinction one would have to specify an unfeasible beam efficiency of well over 99.9% to ensure that AWVR brightness temperature measurement would not be corrupted by more than the 0.01 K stability requirement. A distinction between near and far sidelobes is important in this case because near-sidelobes, which may degrade the beam efficiency, do not necessarily corrupt the WVR given a homogeneous sky brightness temperature. Far-sidelobes are a greater concern since they are subject to much larger temperature swings. The approach taken with the AWVR was to design the antenna as a clear aperture antenna so that far-sidelobe effects could be reasonably predicted by estimates of the main reflector spillover. The AWVR feedhorn design provides a very low edge illumination of less than -25dB. The feedhorn is also designed with a gaussian aperture taper to further reduce sidelobes beyond the reflector edge. Preliminary simulations predicted that not more than 0.3% of the antenna feed power would spill over the edge of the main reflector, that the contribution from backlobes would not exceed 0.5K in the total brightness temperature, and that fluctuations in this power on the 1,000 second time scale of interest to the Cassini GWE would not exceed the 0.01K requirement. Subsequent tests of the completed antenna and radiometers have verified that there is no evidence of antenna backlobes at a detectable threshold of 0.01K during repeated azimuth scans of the AWVR pedestals.

Table 1. Short list of AWVR specifications

Frequencies (GHz)	22.2, 23.8, 31.4
Bandwidth (MHz)	600
NEDT, 1,000-10,000 seconds (K)	0.01
calibration uncertainty (K)	0.5
antenna beamwidth (deg)	1.0
pointing uncertainty (deg)	0.05
1,000 second pointing 'jitter' (deg)	0.005

3. Radiometer design

temperature control

The AWVR 'box' of Figure 1 incorporates a number of design features that relate to the stability of the system. First and foremost is the temperature control. Figure 2 gives a cross section view of the AWVR interior to illustrate the design. The AWVR uses two thermoelectric cooler (TEC) circuits and a double insulated enclosure to achieve very tight temperature control of the RF components. The 'main' TEC is attached to the base plate of the weather-tight enclosure, and serves as the heat exchanger to the outside air. This TEC can heat or cool the instrument under the control of a computer. Fans are used to circulate air from the main TEC throughout the main enclosure and around the interior RF enclosure, as illustrated with arrows in Figure 2. The enclosure incorporates air ducts to guide the air over the top of the interior RF enclosure before returning to the main TEC. The computer, which is embedded in the "electronics" of Figure 2, reads temperatures from thermistors distributed throughout the radiometer. A weighted average of these temperatures is applied to a software control loop to maintain the average temperature surrounding

the interior RF enclosure. The target temperature for this 'main' control loop is 35C. Diurnally changing gradients among the thermistors surrounding the RF enclosure are typically +/- 0.4C. Also, because of the turbulent air flow within the main cavity, random fluctuations of +/- 0.1C are common. These errors are dampened by the insulation of the interior RF enclosure. Within the RF cavity, the design objective is to minimize all heat sources and heat sinks so that RF components can passively settle to a single reference temperature of 35 C. No fans are used inside of the RF cavity because these only add heat and turbulence. All active RF components, which includes amplifiers and noise diodes, are isolated on a plate that is cooled by a smaller secondary TEC to the same 35C. Local gradients among the active components are about 0.1C, but these gradients are constant with time. Insulation is packed around these components to prevent radiant and convective heat flow into the RF cavity. By this arrangement, temperatures within the RF cavity can be controlled and stabilized to +/- 0.004 C.

humidity control

Along with temperature control, humidity control within the radiometer is also needed to maintain stability of the RF circuits. This sensitivity was observed in laboratory test which showed that the forward coupling of a waveguide directional coupler could be changed significantly with a moderate change in humidity. These errors were tied to changes in the dielectric of the air within the waveguide which affected the standing waves along the coupled arm of the waveguide. To control the humidity of the AWVR the enclosure is sealed and a nitrogen purge is applied for at least one day to remove any moisture. It is important to purge the radiometer because moisture will otherwise tend to migrate into the insulation near the exterior walls of the radiometer at night, and then back into the RF cavity when the walls heat up in the daytime. In the field, and if no nitrogen purge is available, air can be vented to the outside via a silica desiccator. Venting is needed to allow for barometric changes.

Radome

The radome of the AWVR provides a moisture and thermal barrier to the antenna feedhorn, as shown in Figures 1 and 2. Expanded polystyrene (EPS) is used for the thermal insulation of the radome in two layers; one on the RF box of Figure 2, and one on the exterior enclosure. To reduce RF reflections all surfaces of the foam are treated with a pattern of v-groove cuts (not visible in the figures). The exterior moisture barrier consists of a thin sheet of Teflon. As shown in cross-section in Figure 2, this Teflon sheet is also modified with an impressed cone shape. This cone shape is needed to disperse reflected power away from the feedhorn aperture.

Dicke switch

Figure 2 also shows the basic layout of the AWVR's waveguide circuits following the horn antenna. The output of the horn connects to WR-34 waveguide to the right of the figure (drawn as a heavy black line) which then feeds a Dicke switch below. The Dicke switch in this case is a unique mechanical voice coil actuated vane attenuator, which is also shown in cross section in the figure. This switch provides the reference black body temperature to the radiometer by periodically inserting a vane attenuator into a broad-wall slotted waveguide. The voice coil actuator moves the vane up and down approximately 0.7cm at a rate of about 3 Hz with a settling time of about 10 ms. The voice coil itself is at the base of the assembly outside of the RF cavity (for cooling) and is connected to the vane by an insulating fiberglass and foam tube. This tube is held by loudspeaker 'spiders' that guide the moving assembly in a straight line without any sliding bearings or contacting parts. The reliability of this switch is excellent: after over two years of continuous operation, or about 200 million switch cycles, there has been no evidence of degradation in either of the two AWVR's.

The RF electrical characteristics of the vane attenuator is one of the keys to the AWVR performance, and was the prime motivation to developing the actuator. This device works over the entire waveguide band, the insertion loss in the 'on' state (i.e. when the vane is pulled out of the slot) is less than 0.3dB, the loss in the 'off' state (vane fully inserted) is more than 45 dB, and in either state of the switch the return loss is well in excess of 30dB. To date, such performance can not be matched by any known electronic switch.

The loss of this device in the 'on' state is also very reliable and stable since, in this state, it essentially is a straight piece of waveguide.

Noise diodes

Figure 2 show a set of three cross-guide couplers following the Dicke switch. Each of the couplers connects via a coaxial cable to a corresponding noise diode on the TEC plate below. The noise diodes are stabilized in temperature and are carefully controlled in terms of bias current and duty cycle to provide a repeatable and stable additive noise reference to the radiometer. These noise diodes are switched on and off via their bias currents under the control of the computer, and the responses are used monitor and stabilize the gain of the receiver. The effective noise temperatures of the noise diodes, as referenced to an equivalent source at the antenna, is nominally 300 ~ 500 K in the two AWVR's. The stability with which the output of the noise diodes can be compared to power at the antenna input is a critical determinant of the radiometer stability. The fact that there are three noise diodes is another key to the AWVR performance.

One of the most difficult problems concerning the use of coupled noise diodes is the multitude of small reflections and associated standing waves that are inevitable in the interconnecting transmission lines of the radiometer front-end. The laboratory work preceding the AWVR development revealed just how problematic standing waves are, and how difficult it is to compare the power coupled at one point in a transmission line with the power coupled elsewhere along the same line. The problem is that a small change in a reflection anywhere between the radome, the receiver, and the noise diode can coherently interfere with the coupler to cause a large change in the effective noise diode temperature. Such problems are well known in the microwave industry, but they are too often overlooked in radiometer designs. A simple innovation of the AWVR in this regard is the use of redundant noise diodes. By injecting several noise diodes at different points along the transmission line the system can at least measure- and to a degree mitigate- variations in the noise diode coupling. This feature is invaluable as a diagnostic tool as it provides a continuous and relevant measure of the system performance. Relative changes in the noise diode outputs are continuously monitored in the data system. On many occasions these changes have been used to quickly identify and repair some very subtle problems in the AWVR. For example, internal humidity changes, movement of the radome, and a variety of coaxial cable problems have all been identified with this technique. The fact that the noise diodes are injected after the Dicke switch also helps by providing such comparisons in either state of the Dicke switch.

Receiver

The output of the couplers in Figure 2 is connected to a ferrite isolator, which in turn passes the signal through a waveguide-to-coaxial adapter, and then to the low noise amplifier (LNA) on the RF TEC plate below. The isolator prevents the relatively poor mismatch and the cold noise temperature of the LNA from affecting the calibration circuits above. The LNA has a noise figure of approximately 3dB from 22GHz to 32GHz. Following the LNA, the remainder of the receiver (not shown in the figure) is a 3-channel single-sideband direct-detect system: a frequency triplexer divides the RF signal into the three observing bands, which are again amplified, and then passed to three zero-bias Shottky detectors. All of these devices are commercial coaxial components, and all are attached to the same TEC plate at the base of the RF enclosure. The RF power level reaching each of the detectors is approximately -17dBm. OP-27 video amplifiers then boost the detected signal to about 2 volts, which is then passed to digitizer/integrators in the electronics assembly below the RF enclosure. The digitizer/integrators consist of voltage-to-frequency converters followed by digital counters, as is common in radiometer design.

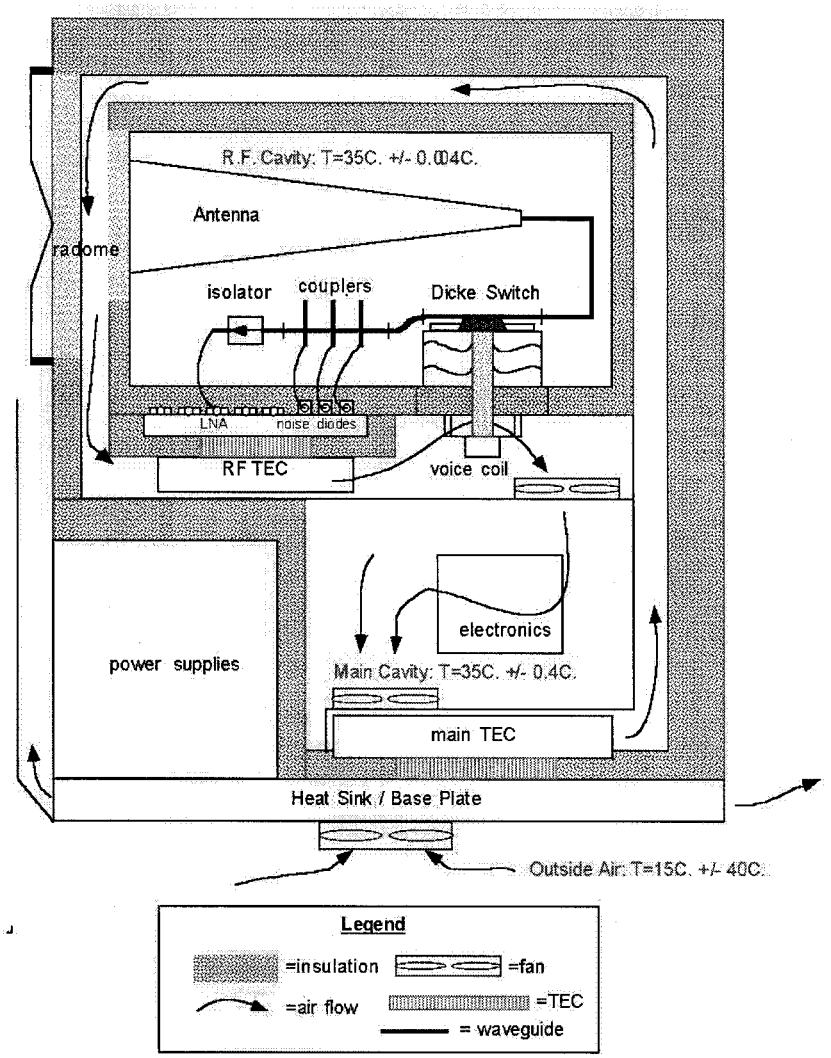

Figure 2- AWVR cross-section

Timing

The operation of the Dicke switch and each of the noise diodes follows the timing diagram of Figure 3. A simulated video waveform is also illustrated in Figure 3. The timing of the measurement modes is determined by speed limitations of the mechanical Dicke switch, and a desire to operate the noise diodes at a constant duty cycle and with very short pulses in order to minimize thermal cycling of the devices. The sequence of measurements is controlled by an interrupt service routine in the computer. Digital hardware generates interrupts based on a 10 MHz crystal every 3.2768 milliseconds, and the interrupt service routine then reads the digitizer/integrators and programs the next calibration mode on each measurement cycle. The digitized data is then buffered by the interrupt routine, and then read by a data acquisition program which sorts and averages the data by frequency and calibration mode. In the end, the data system records 24 values representing 3 channels (22.2, 23.8, and 31.4 GHz) in 2 modes of the Dicke switch (antenna or reference), and 4 noise diode modes (noise diodes A, B, or C, and no-noise-diode).

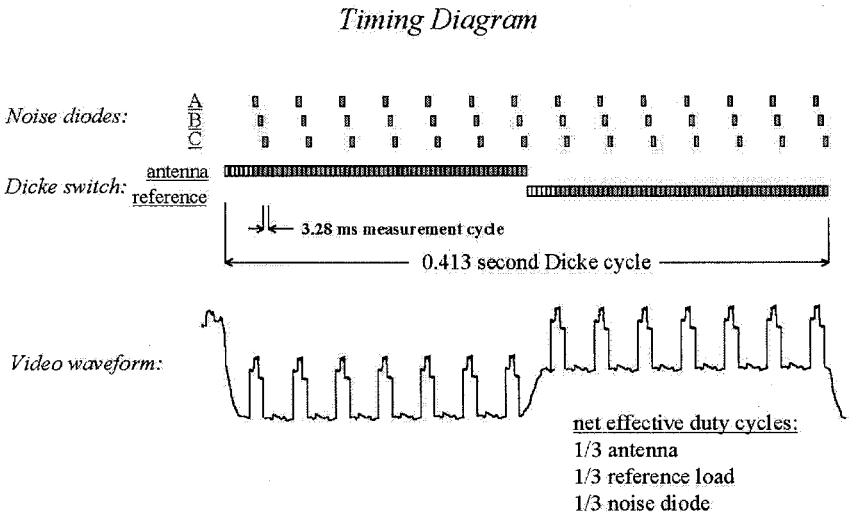

Figure 3- Timing diagram

4. Calibration

The calibration of the AWVR involves two steps: (1) a system calibration to determine the noise diode temperatures, and then (2) application of these noise diodes temperatures to the routine radiometer observations. The system calibration (1) occurs on a once-per-month basis, and involves a tip-curve procedure as discussed below. In step (2) the noise diodes are used to transfer the system calibration to the routine observations according to

$$T_B \cong T_o - \frac{C_o - C_s}{C_{NDD}} T_{ND} \tag{1}$$

where T_B is brightness temperature, T_o is the physical temperature of the Dicke switch reference load, C_o is the measured "counts" in response to this load (the word "counts" here is a reference to the counter of the digitizer, and is synonymous with "integrated detector voltage"), C_s is the "sky" counts as measured during the antenna mode of the Dicke switch, C_{NDD} is a sum of the counts-deflected for all the modes of the noise diodes, as defined below, and T_{ND} is a sum of the corresponding noise diode temperatures. There are three noise diodes which are injected in the two modes of the Dicke switch for a total of six noise diode measurements. The noise diode terms of (1) expand to:

$$C_{NDD} = C_{sNDA} + C_{sNDB} + C_{sNDC} - 3C_s + C_{oNDA} + C_{oNDB} + C_{oNDC} - 3C_o \tag{2}$$

and

$$T_{ND} = T_{sNDA} + T_{sNDB} + T_{sNDC} + T_{oNDA} + T_{oNDB} + T_{oNDC} \tag{3}$$

where the 's' and 'o' subscripts denote sky and reference modes, respectively, the "A", "B", and "C" subscripts refer to each of three noise diodes, and C_o and C_s are the same reference and sky counts from Equation 1.

In equation 1, T_o is readily measured with a thermistor placed next to the absorbing vane of the Dicke switch, and is very close to 308.14 K (=35 C). This leaves only T_{ND} in Equation 1 to be calibrated by the tip-curve procedure. A good description of this procedure can be found in [1]. Briefly, this procedure involves observing the sky at various elevations, and then extrapolating the measured curve of observed counts versus calculated air mass (= secant of elevation assuming a horizontally uniform atmosphere) to zero air mass. At zero air mass the extrapolated counts correspond to the well known microwave cosmic background temperature of 2.7K, and the needed cold target reference is established. To a first order this extrapolation is liner, but in practice a more accurate extrapolation involves an iterative solution to account for the non-linear relation between brightness temperature and atmospheric opacity. The accuracy of the procedure is limited largely by the degree to which the atmosphere is horizontally uniform, and by uncertainties in the temperature of the atmosphere. Most of the uncertainties scale with the humidity of the atmosphere, so the quality of the tip-curve tends to be best on dry days. Typically, a single elevation scan- from near the horizon to zenith- is sufficient for about 1 K of uncertainty in the reference. Over the course of a day the averaging of multiple tip curves can reduce the instabilities to about 0.1K. To build up the best possible calibration data base fore the AWVR's, the tip-curve procedure is performed continuously any time that the radiometers are not required for other observations.

5. Performance Evaluation

The following performance evaluation is broken into three parts: (1) "delta-T" to compare the short term noise of the AWVR with theoretical limits; (2) "stability estimates" to establish the performance on time scales of about 1,000 seconds to 2-days using Structure Function and Allen Standard Deviation analysis; and (3) "tip-curve" results which track the noise diode calibration history for nearly a year of operations at Goldstone.

delta-T (thermal noise)

The theoretical noise limit, or 'delta-T', of the AWVR can be computed from Equation 1 given the duty cycles of Figure 3. The formulation of the delta-T equation is straight forward, but quite lengthy, and leads to the following formula:

$$\Delta T = \sqrt{\frac{3\left(1+3\frac{T_o - T_B}{T_{ND}}\right)^2 (T_r + T_o)^2 + 3\left(1-3\frac{T_o - T_B}{T_{ND}}\right)^2 (T_r + T_B)^2 + 18\left(\frac{T_o - T_B}{T_{ND}}\right)^2 \sum_{x,y}(T_r + T_x + T_{xNDy})^2}{B\tau}} \quad (4)$$

where B is the detection bandwidth of the AWVR (about 600 MHz), τ is integration time, T_r is the receiver noise temperature (about 500K), x refers to the Dicke switch mode, y refers to a specific noise diode, and the remaining terms are from Equation 1 and 3.

To measure the thermal noise of the two AWVR's the radiometers were operated side-by side while observing the sky. The difference between the brightness temperatures of the two radiometers at comparable frequencies was then used to estimate the combined delta-T according to:

$$E_S(\tau) \equiv \sqrt{\frac{1}{2}\left\langle \left(D(t,\tau) - D(t+\tau,\tau)\right)^2 \right\rangle} \quad (5)$$

where D is the brightness temperature difference as a function of time, t, and integration time, τ, and <> is the expectation operator. Note that E_S equals the standard deviation of D when D is zero-mean white noise.

Figures 4 summarize the combined delta-T of the A-1 and A-2 radiometers as computed from measurements with Equation 5 (solid lines) and from the root-sum-square (RSS) of the theoretical estimates of Equation 4 for the two radiometers (dashed lines). Figure 4 is a log-log plot of noise versus integration time, starting at a minimum integration set by the Dicke switch cycle time of 0.413 seconds. As can be seen in this figure, the radiometer noise agrees very well with the theory out to an integration time of 1,000 seconds, at which point instability of the two radiometers prevents any further noise reduction. The minima near 1,000 seconds of the three channels is approximately 0.005 to 0.008K. Since Figure 4 represents the difference of two nearly identical radiometers, these numbers can be divided by SQRT(2) to estimate that noise of individual radiometers is approximately 0.0035 to 0.0056K. At 1,000 second time scales the Cassini requirements have been translated to a 10mK radiometric stability requirement. From Figure 4, the AWVR's will evidently meet this requirement comfortably.

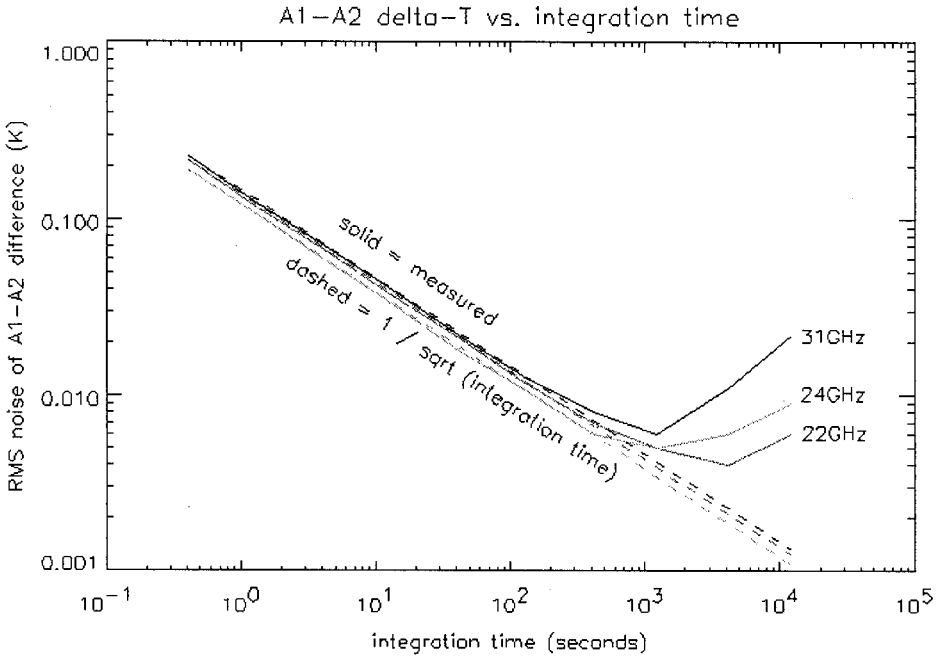

Figure 4- Theoretical (dashed) and measured (solid) delta-T from side-by-side comparisons of the two AWVR's

Stability Estimates from Side-by-Side Data

The measured delta-T's presented above were calculated from the difference of consecutive samples separated by the integration time of the box-car averages. This calculation is useful for determining where the instrument stability exceeds thermal noise, but it is not necessarily the best measure of instrument stability at longer time scales since integration times exceeding 1,000 seconds are of little practical interest. Better measures of the instrument stability on long time scales are provide by either structure function analysis or by the Allen Standard Deviation of the A1 - A2 brightness temperature difference. As a sample case for the following analysis, Figures 5 and 6 give the brightness temperature and brightness temperature differences, respectively, from sky measurements spanning 5 days. A box-car integration time of 124 seconds applies. In Figure 5 the overlay of brightness temperatures of the two radiometers match too well to discriminate between the A1 and A2 results, but in Figure 6 there is an obvious drift in the 31 GHz channel that will have some bearing on the following error analysis.

Figures 7 and 8 give two different estimates of the AWVR "excess" errors- meaning errors in excess of the inherent thermal noise- versus time scale. Figure 7 was computed as the square root of the structure function after subtracting the thermal noise according to

$$E_S(\tau) \equiv \sqrt{\left\langle (D(t) - D(t+\tau))^2 \right\rangle - \left\langle (D(t) - D(t+\tau_0))^2 \right\rangle} \qquad (6)$$

where D is the A1 - A2 brightness temperature difference from Figure 6, t is time, τ is the time difference for the horizontal scale of Figure 7, and τ_0 is the integration time. τ_0 is set sufficiently small so that the right-most term of Equation 6 is an estimate of the thermal noise only. In Figures 5 through 8 τ_0 is 124 seconds where, as evident in Figure 4, the noise is consistent with the theoretical thermal limit. Subtracting thermal noise in this manner isolates errors caused by slower instrumental instabilities which are largely independent of integration time for $\tau_0 < \sim 1000$ seconds. Figure 7 is thus applicable to all integration times of interest. As can be seen in Figure 7 the 22 and 24 GHz channels both exhibit an excellent stability of

less than 0.03 K excess structure error. The minima near 86,000 and 172,000 seconds- or 1 and 2 days- in the 22 and 24 GHz channels also indicate that much of the error is diurnal. The large drift in the 31Ghz channel, in the other hand, obviously contributes a great deal to the excess structure error of Figure 7.

The original requirements for Cassini GWE are expressed as an Allen Standard Deviation (ASD) of the path delay correction. ASD is insensitive to linear drift, so the 31 GHz drift of Figure 6 should be less of a concern for the GWE than indicated by the structure function. To evaluate ASD errors on the same scale with Figure 7, Figure 8 gives a modified version of the Allen Standard Deviation (ASD) in which the units are Kelvin- instead of the conventional seconds/second ASD definition- and the thermal noise has again been subtracted according to

$$E_A(\tau) \equiv \sqrt{\frac{1}{2}\left\langle \left(D(t) - 2D(t+\tau) + D(t+2\tau)\right)^2 \right\rangle - \frac{1}{2}\left\langle \left(D(t) - 2D(t+\tau_0) + D(t+2\tau_0)\right)^2 \right\rangle} \qquad (7)$$

Note that there is no division by τ in Equation 7, as would be convention in the ASD calculation.

Figure 8 exhibits a clear diurnal component in all three channels. The 31 GHz channel of Figure 8 again exhibits the worst stability of the three AWVR channels, but the minima of about 20 or 30 mK at 1 and 2 days time lag are remarkably small and consistent.

The 31GHz performance is generally poorer than the 22 and 24 GHz channels due to known increases in the reflection coefficients of the radome and in nearly all RF components of the radiometer's front-end. Poor directivity in the directional couplers used for the noise injection at 31 GHz are also problematic. The particular data set of Figure 5 was collected shortly after a modification to one of the A-1 radiometer's noise diodes and a coaxial cable, and the drift in the 31GHz channel can be linked to a comparable drift among the different noise diode data of that channel. This drift has since stopped.

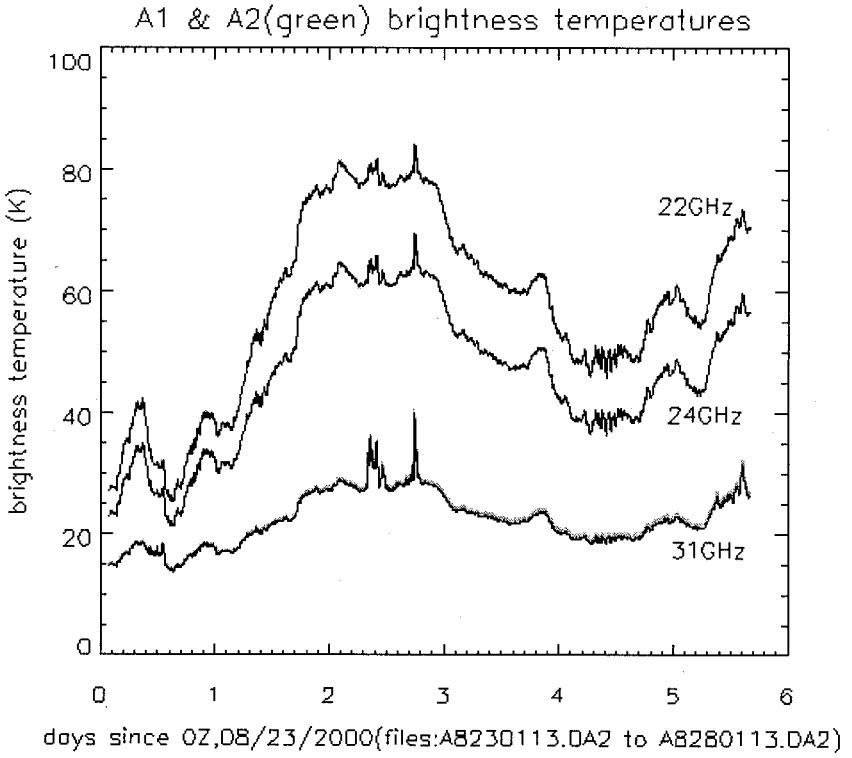

Figure 5- Brightness temperature overlay of the A1 radiometer (black) and A2 (green). Only the 31 GHz channels of A1 and A2 differ enough to discriminate between the two instruments.

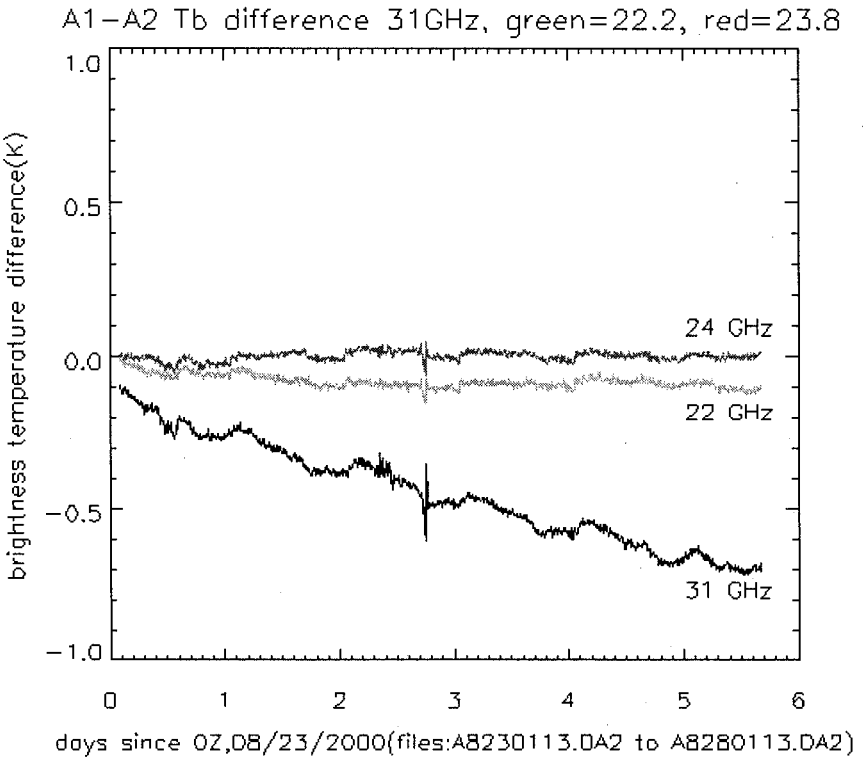

Figure 6: Difference of A1 and A2 brightness temperatures from Figure 5. Integration time = 124 seconds.

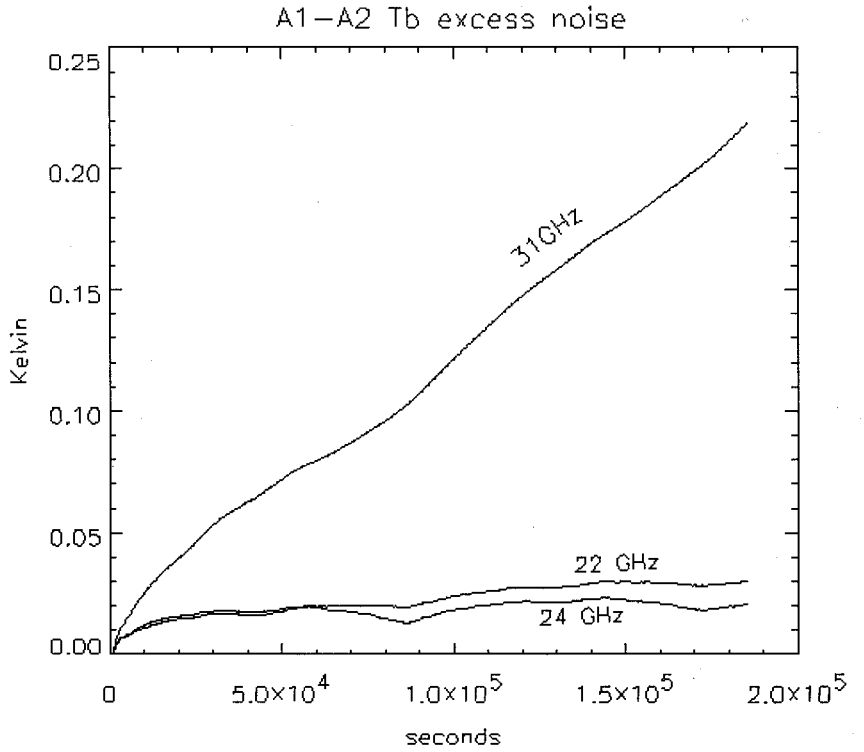

Figure 7- Excess noise of A1-A2 brightness temperatures difference of Figure 6 computed as SQRT(structure function - thermal noise).

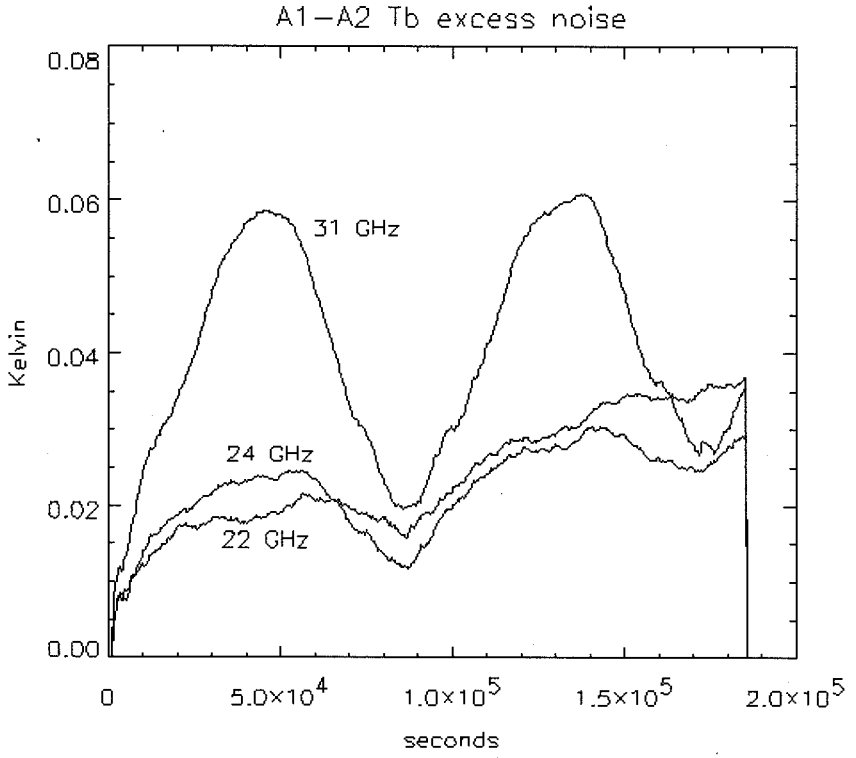

Figure 8- Excess noise of A1-A2 brightness temperatures difference from Figure 6 computed as a pseudo-Allen Standard Deviation by Equation 7.

AWVR stability estimates from tip-curve data

The long term stability of each of the AWVR's can be independently estimated from tip-curve calibration data collected over the past year at Goldstone. Figures 9 and 10 summarize errors estimated by subtracting brightness temperatures based on a one-time calibration from those which have been recalibrated from daily tip-curves. As noted above, the primary result of a tip curve calibration is an estimate of the effective noise diode temperature, so Figures 9 and 10 essentially track the histories of the noise diodes with respect to a one-year average. Only high quality tip curve results which span an entire day are plotted here; out of a possible 300 days, only about 90 days were selected. Most days were rejected due to poor weather or outages caused by other experiments or maintenance. As can be seen, both AWVR's have drifted very little over nearly one year. Only minor long term errors in the 31 GHz channels are evident. Otherwise, much of the short-term scatter- especially in the 22 GHz channels- can be attributed inhomogenaities in the atmosphere which limit the accuracy of the tip-curve result. Table 2 summarized the standard deviation of Figures 9 and 10. Also given in Table 2 are standard deviations calculated from the 'best' month in which particularly dry and stable weather yielded very high quality tip-curve calibrations.

The last three rows in Table 2, and Figures 11 and 12, provide estimates of the average calibration errors versus time of day. These were calculated by binning the tip curve data in 6-hour segments sorted by time of day, and then comparing the results with the all-day average. These calculations were repeated for various sub-sets of the tip curves (spanning different seasons) and the results were consistent with the averages of Figures 11 and 12. Thus, Figures 11 and 12 are fairly accurate estimates of persistent diurnal errors of the AWVR's. The standard deviation of the difference of A1 and A2 diurnal errors, as presented in the last row of Table 2, are also in good agreement with the errors that were estimated by the ASD of Figure 8 using the side-by-side AWVR comparisons.

Known electrical losses in the AWVR antenna and radome may account for much of the diurnal errors of Figures 11 and 12. Radiometric tests of flat reflector samples were made prior to antenna fabrication which indicated that reflector losses would contribute 0.5 to 0.9K to the antenna noise temperature. Add to this the estimated reflector spillover of 0.5 K, and the total antenna noise contribution from elements outside of the WVR's thermal enclosure will likely exceed 1.0 ~ 1.5K. These values scale to approximately 0.03K to 0.05 K of peak-to-peak antenna noise temperature change given an average day-to-night ambient temperature change of 10K out of an average 300K. Such noise temperature changes apparently agree with Figures 11 and 12.

Table 2. Brightness temperature errors calculated from day-to-day tip-curve results of Figures 9 and 11.

RMS TB errors from daily tip-curve vs. noise diode calibrations (K):

	31GHz	22GHz	24GHz
~ year (RMS of Figures 9 and 10):			
A1, 94 days between 7/20/1999 and 6/5/2000:	0.204	0.162	0.151
A2, 89 days between 7/20/1999 and 6/6/2000:	0.431	0.185	0.093
~ best month:			
A1, 22 days from 12/22/1999 to 1/20/2000:	0.086	0.057	0.041
A2, 29 days from 12/10/1999 to 1/15/2000:	0.027	0.064	0.045
~ diurnal errors from Figures 11 and 12:			
A1 vs time of day:	0.031	0.037	0.027
A2 vs time of day:	0.057	0.039	0.031
difference of A1-A2 vs. time of day:	0.072	0.035	0.022

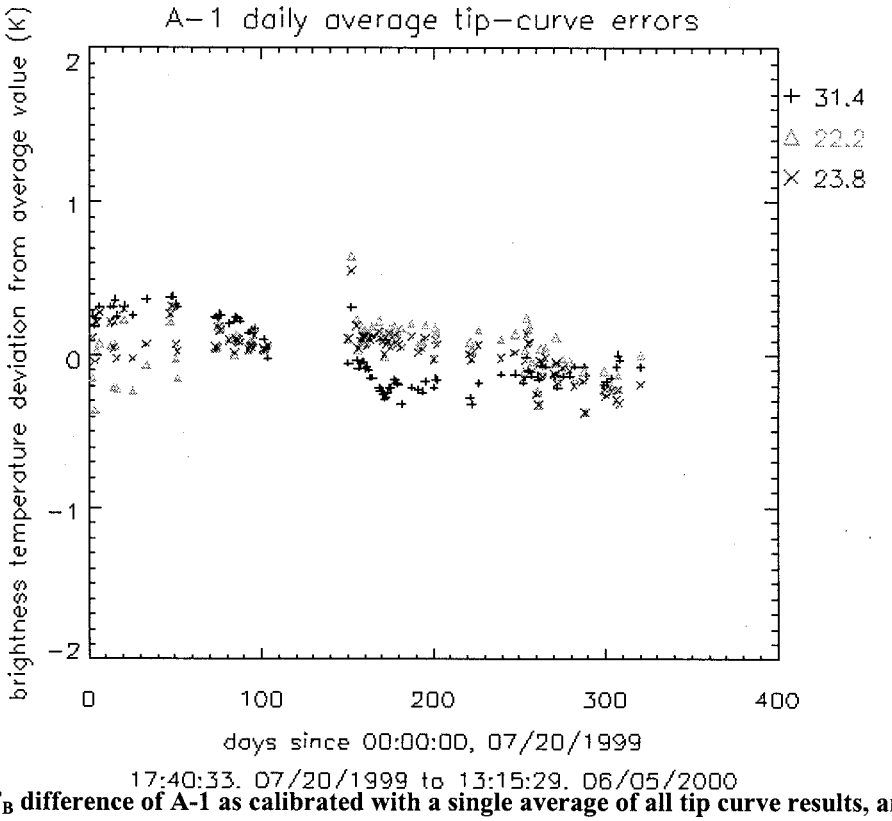

Figure 9: T_B difference of A-1 as calibrated with a single average of all tip curve results, and calibrations based on daily tip curve results. Positive error means that the long-term calibrated AWVR overestimated the brightness temperature on a given day.

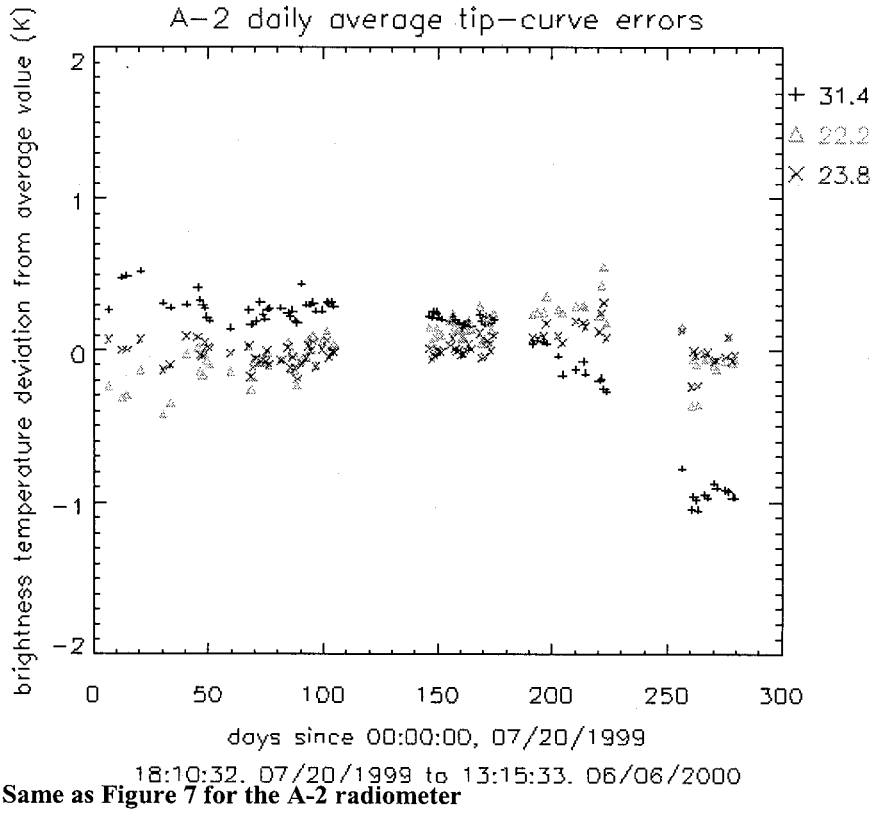

Figure 10- Same as Figure 7 for the A-2 radiometer

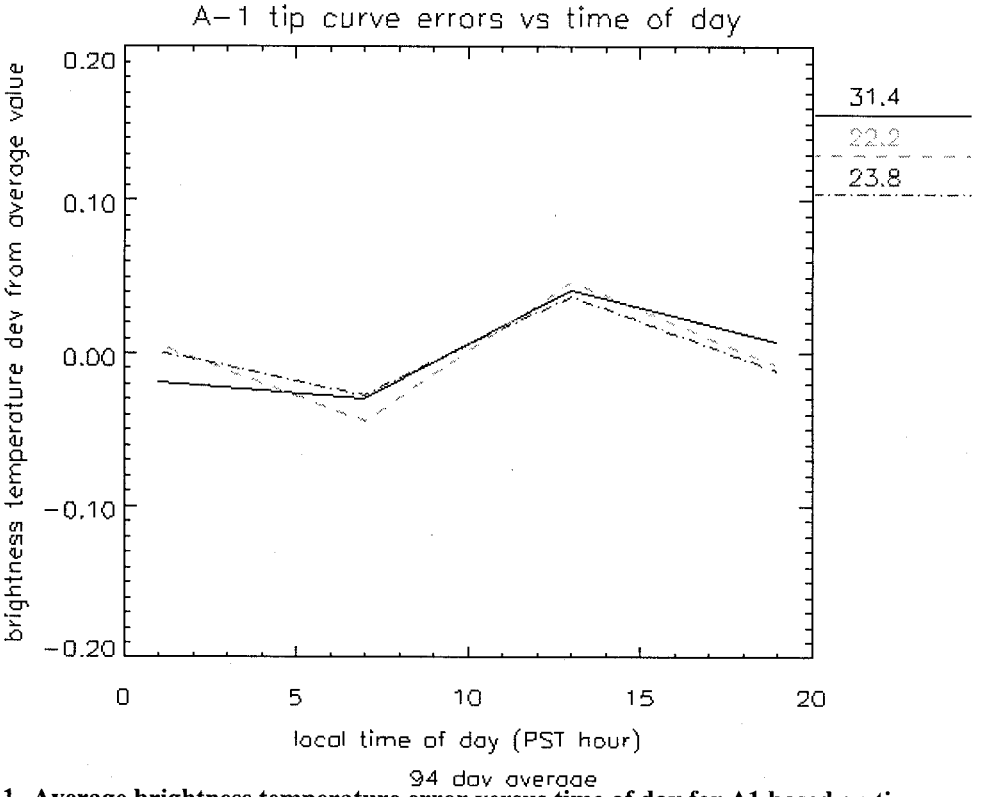

Figure 11- Average brightness temperature error versus time of day for A1 based on tip curve analysis

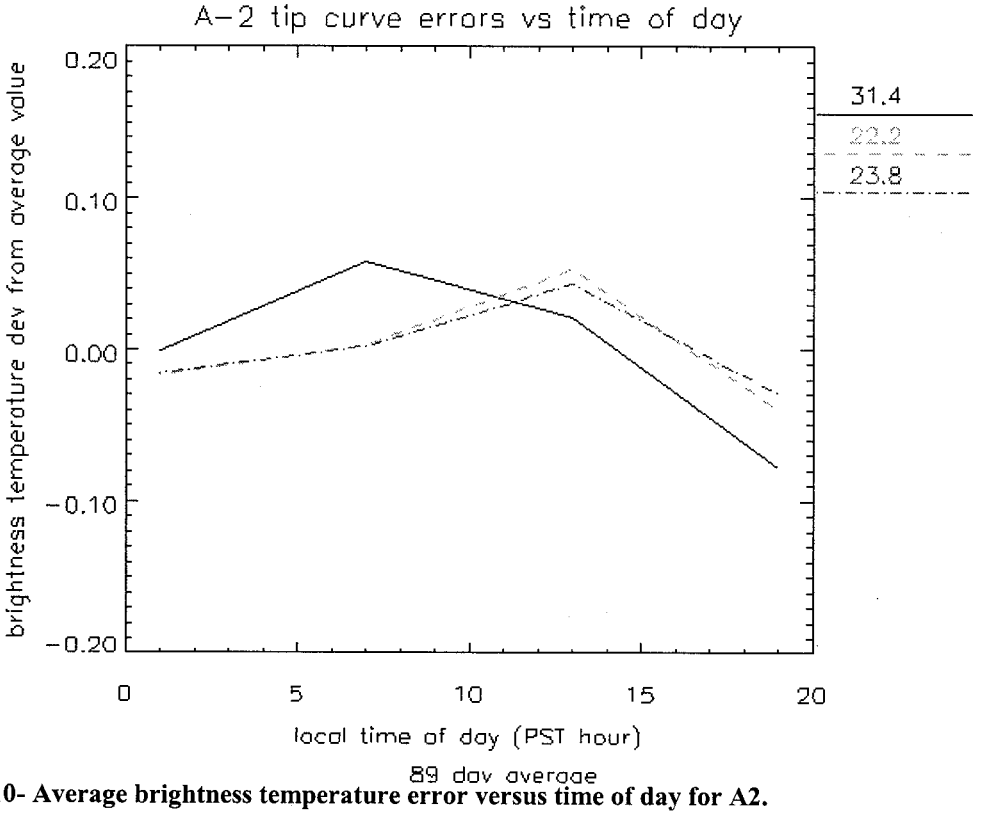

Figure 10- Average brightness temperature error versus time of day for A2.

6. Conclusion

The AWVR's will meet the stability requirements of the Cassini GWE experiment. On time scales of 1,000 to 10,000 seconds the above data demonstrates that the 0.01K stability goal has been met, with some channels approaching 0.003K. For a radiometer in which the observed brightness temperature is some 300 K colder than the internal (Dicke) reference temperature, a 0.003K stability figure implies a gain stability of 10 parts per million. Such precision would be considered good if the measurement were distance, weight, or voltage- but for a microwave power measurement such precision is outstanding. Beyond the Cassini requirements- and well past 10,000 second time scales- the AWVR's also exhibit an excellent level of stability of about 0.02K for day-to-day errors, and approximately 0.03 ~ 0.08 K per month.

This report has focused on the stability of the radiometers, but not the absolute uncertainty. In Table 1 the quoted calibration uncertainty was given as 0.5K. Yet no data has been provided here to support this estimate. 0.5K is, at present, a best estimated based on comparisons with other WVR's and on a number of factors which are beyond the scope of this paper. These factors include detector non-linearities, antenna losses that vary with elevation angle, and a host of potential errors associated with the tip curve procedure (e.g. earth curvature, refraction, finite antenna beamwidth, etc. [1]). A more careful error analysis is warranted for future work.

Acknowledgements. The work described was performed at the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

References

[1] M.A.Janssen, *Atmospheric remote sensing by microwave radiometry*. New York: Wiley & sons, 1993.

[2] S.J. Keihm and K.A. Marsh, "New model-based Bayesian inversion algorithm for the retrieval of wet troposphere path delay from radiometric measurements," *Radio Sci.*, vol. 33, no. 2, pp411-419, March-April 1998.

[3] R.P. Linfield et al., "A test of water vapor radiometer based troposphere calibration using very long baseline interferometry observations on a 21-km baseline", *Radio Sci.*, vol.31, pp. 129-146, 1996.

[4] M. Tinto and J. W. Armstrong, "Spacecraft Doppler tracking as a narrow-band detector of gravitational radiation," *Physical Review D.*, vol.58, pp. 042002-1 -8, July 1998.

[5] A. B. Tanner, "Development of a high stability water vapor radiometer," *Radio Sci.*, vol. 33, no. 2, pp. 449-462, Mar.-Apr. 1998.

[6] R.P Linfield and J.Z. Wilcox, "Radio metric error due to mismatch and offset between a DSN antenna beam and the beam of a tropospheric calibration instrument," TDA Progree Report, 42-114, pp.1-9, 1993.

[7] R. P. Linfield, "Effect of aperture averaging upon tropospheric phase fluctuations seen with a radio antenna," *Radio Science*, vol. 33,. no. 5, pp. 1353-1359, Sept.-Oct. 1998